A Modular Digital Front End Trigger System for the DØ Experiment

Fermilab

Background

Fermilab is currently home to the world's most powerful particle accelerator, the Tevatron. This machine accelerates bunches of protons and anti-protons to 99.999% the speed of light and then collides them head on in the center of massive detectors. When these collisions occur, sub-atomic particles are produced and travel outward and penetrate layers of sensitive detectors which characterize the particles in terms of energy, flight path, etc. Recently, Fermilab upgraded both of the colliding beam detectors (CDF and DØ) with new hardware. These improvements were necessary to handle the higher collision rates that these detectors would see when the upgraded accelerator resumed operation.

At high collision rates data is produced much faster than it can be analyzed in real-time, and thus hardware must be used to quickly determine if a collision has produced any useful physics data. To implement this DØ has developed a three tiered trigger system. The first tier, called Level-1 (L1), analyzes the detector data in real-time and quickly makes a decision as to whether or not the collision was "interesting". If so, the L1 hardware informs the second tier L2 hardware and a more detailed readout of the data occurs. The last tier consists of the L3 hardware, which dumps out all of the raw data for off-line analysis and debugging.

The Problem

In the design phase it quickly became apparent that even the fastest processors could not handle the extremely high data rates that the L1 trigger hardware would have to analyze. The decision was made to use Field Programmable Gate Arrays (FPGAs) to crunch on the data in parallel. Several crates of FPGAs operating in parallel would be needed to handle aggregate data rates that would exceed 1 terabit per second.

A Solution

An efficient hardware trigger system would be one that could adapt to diverse trigger configurations while keeping the number of unique board designs to a minimum. The Digital Front End (DFE) boards were developed around this idea – a modular system of motherboards, daughterboards, and transition boards that form flexible building blocks that can be assembled in different ways to meet the requirements of many diverse detectors.

DFE Motherboard

The Motherboard is a custom 6U x 320mm board which accepts ten 1.85 Gbps Low Voltage Differential Signal (LVDS) channel links. These high speed serial links are converted back to a parallel bus and buffered before driving the signals up to the daughterboard connectors. The motherboard handles all communication with the DFE backplane, implements board level diagnostics, and provides a data path for the four 28-bit output busses. A single DFE motherboard design is used for all DFE implementations.

DFE Daughterboards

The purpose of the daughterboard is to support the FPGAs. Two different daughterboard designs were produced: a "single wide" model contains five Xilinx Virtex FPGAs, and a "double-wide" model which has three large Virtex FPGAs. The "single wide" daughterboard can see only seven out of the ten LVDS motherboard inputs, while each FPGA on the "double-wide" daughterboard can see all ten LVDS motherboard inputs.

The firmware residing in the FPGAs varies substantially depending on the application. However, the top level block diagrams tend to show some similarities – an example is shown to the right. Before the data can be processed it must be synchronized with the DFE motherboard clock (link de-skew module). The motherboard interface and L3 sender modules are for diagnostics, while most of the application-specific work gets done in the core module.

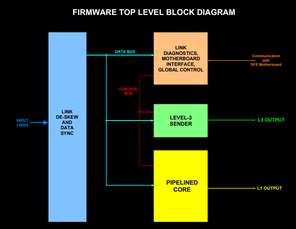

DFE Transition Board

The transition board is plugged into the back of the DFE crate. It allows the four output busses to drive several different types of media: LVDS Channel Links, serial co-axial copper cables, and fiber optics. The serial co-ax and fiber optic drivers are contained on small daughterboards that plug into the transition board. Small FIFO memories are used to remove all clock jitter from the transition board's input busses, thus providing clean clocks and data to the sensitive gigabit serializer daughterboards.

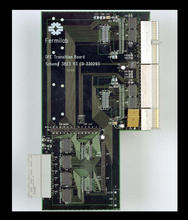

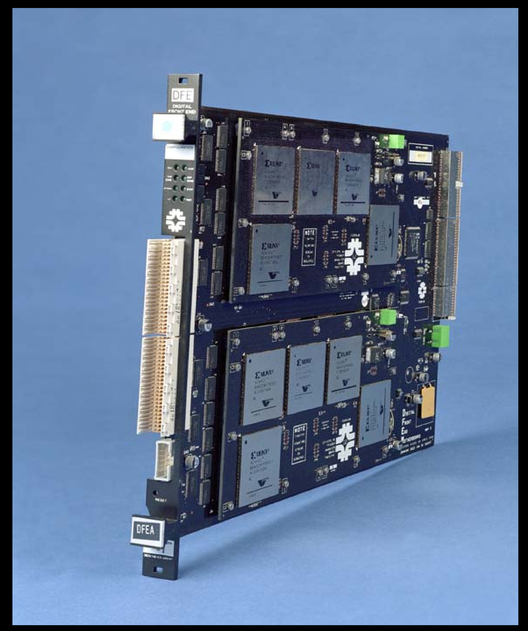

DFE Motherboard with two "Single-Wide" Daughterboards.

DFE Crate Controller

The primary task of the DFE Crate Controller is to handle low-speed communications between the DFE Motherboards and the rest of the DØ experiment. The MIL-STD1553B protocol was selected for this communications bus. The 1553 bus is reliable and very low-noise, but it's also very slow, especially when the hundreds of FPGAs require megabytes of configuration data – which would require hours to download.

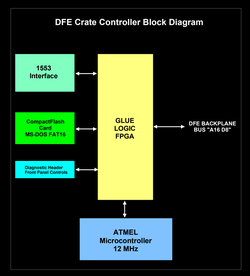

The solution was to develop a crate controller with enough memory on it so that it could hold all of the FPGA configuration files locally. Then, when the FPGAs needed to be configured this could be done quickly just by issuing a few 1553 commands. CompactFlash cards (like those used in digital cameras) were selected for their reliability and low cost. The CompactFlash card is formatted as an MS-DOS hard disk, which means that it can be removed from the crate controller and inserted into a laptop PC. This is useful for quickly downloading files or checking file integrity.

A small 12MHz RISC microcontroller forms the heart of the DFE Crate Controller. The software running on this microcontroller is basically a list processor. A list of commands are written into a buffer and then the microcontroller is told to execute the commands. Some command examples are: reset a DFE board, delete a file, calculate file checksum, append data to file, and configure DFE board with file *x*, etc. When the crate controller is not executing commands it collects status information from the DFE boards in the crate.

Design Application: Central Fiber Tracker

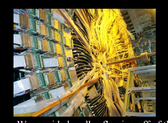

Waveguide bundles flowing off of the Central Fiber Tracker

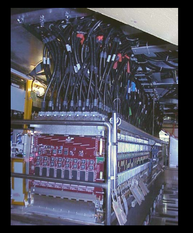

Analog Front End Boards on Cryostat, waveguide bundles drop down from the Central Fiber Tracker

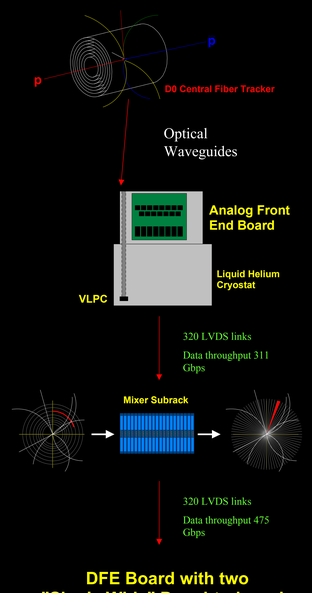

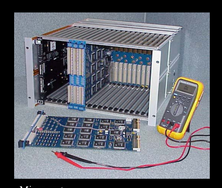

Mixer crate

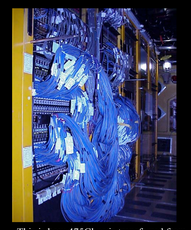

This is how 475Gbps is transferred from one relay rack to another…

LVDS Cables entering the last two tiers of DFE boards.

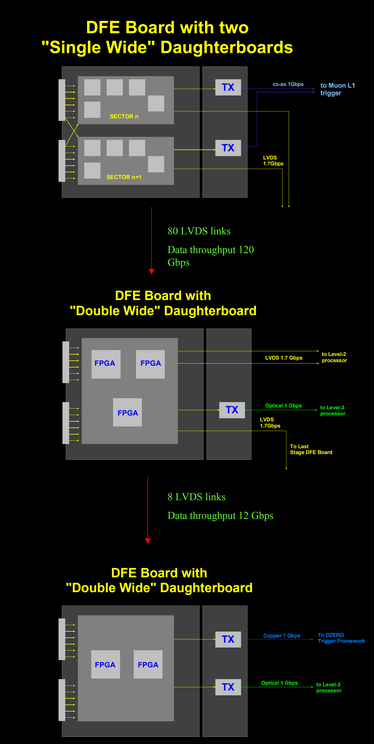

Protons and Anti-Protons collide in the center of the Fiber Tracker every 400 nanoseconds. This is called a "crossing" or "event". As particles travel outward from the collision point they pass through concentric cylinders of scintillating fibers. As the particles travel through the scintillating fibers they create photons.

Over 38,000 optical waveguides deliver the light from the Central Fiber Tracker to the Cryostat / AFE board assembly. The cryostat contains very sensitive photo-diodes called Visible Light Photon Counters, which are cooled to 9K. At this temperature the diode can detect single photons and convert that energy into a very small electric charge.

The Analog Front End boards see this charge and send it to a discriminator and also to an 8-bit Analog to Digital Converter. The discriminator bits are sent over LVDS channel links to the Mixer crate.

The Analog Front End boards see fiber bits arranged in layers while the tracking algorithms need the fiber data arranged in eighty 4.5 degree sectors. The Mixer crate handles all of this data reorganization with a latency of less than 150ns.

Now that the data is packed correctly it's sent to the first tier of Digital Front End boards. This first tier consists of 40 motherboards, each motherboard having two "single-wide" daughterboards. These 80 daughterboards unpack the fiber data and compare the sector bitmap pattern against 1.2 million potential track patterns. The straightest tracks have the highest transverse momentum (Pt) and are currently the most interesting events from a physics perspective.

Each daughterboard finds up to 24 tracks in each sector. The 6 highest Pt tracks in each sector are sent over gigabit co-axial cables to the L1 Muon trigger, where the tracks are matched to hits in the Muon detector. In addition, counts of the found tracks are passed downstream to the next DFE board. Other data such as sector occupancy (number of fibers hit) is also calculated and sent out at this time.

When this DFE board receives an L1_ACCEPT control signal it switches into L2 readout mode. The DFE board jumps back a fixed number of events and extracts the list of found tracks. This list is zero suppressed and sorted before sending on to the next board.

The next tier of DFE boards concatenates and sorts data within an octant (10 sectors). In L1 mode this board receives the track counts and sector occupancy numbers. It must sum up the number of tracks, determine which sector had the most fibers hit, and check for isolated tracks. In L2 mode the upstream boards send up to 240 tracks, which must be sorted in order of decreasing Pt. This sorted list of tracks is then passed on to a processor crate for further analysis. Switching over to L2 mode also causes this DFE board to dump out a copy of its inputs and send it to the L3 processor for debugging and verification studies.

The last tier consists of a single DFE board which accepts data from all eight octants. Based on sector occupancy and track counts it sets various trigger term bits, which are passed on to the DØ trigger framework. The trigger framework then tells the rest of the detectors to read out a particular event or discard that data.

Overall latency of this trigger chain is 2.5μS.

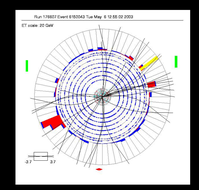

Latest Results:

Triggering and taking data!